arXiv:1608.08505v1 [cs.DS] 30 Aug 2016

Placing Arrows in Directed Graph Drawings*

Carla Binucci[1], Markus Chimani[2], Walter Didimo[1],
Giuseppe Liotta[1], Fabrizio Montecchiani[1]

[1] Università degli Studi di Perugia, Italy
{carla.binucci,walter.didimo,giuseppe.liotta,fabrizio.montecchiani}@unipg.it
[2] Osnabrück University, Germany, markus.chimani@uni-osnabrueck.de

Abstract. We consider the problem of placing arrow heads in directed graph drawings without them overlapping other drawn objects. This gives drawings where edge directions can be deduced unambiguously. We show hardness of the problem, present exact and heuristic algorithms, and report on a practical study.

1 Introduction

The default way of drawing a directed edge is to draw it as a line with an arrow head at its target. While there also exist other models (placing arrows at the middle, drawing edges in a "tapered" fashion, etc.; cf. [6,7]) the former is prevailing in virtually all software systems. However, this simple model becomes problematic when several edges attach to a vertex on a similar trajectory: it may be hard to see whether a specific edge is in- or outgoing, cf. Fig. 1 and Figs. 5 and 6 in the appendix.

We try to solve this issue by looking for a placement of the arrow heads such that (a) they do not overlap other edges or arrow heads, and (b) still retain the property of being at—or at least close to—the target vertices of the edges. In the following, we show NP-hardness of the problem, propose exact and heuristic algorithms for its discretized variant, and evaluate their practical performance in a brief exploratory study. We remark that our problem is related to map labeling and in particular to edge labeling problems [3,4,8,9,10,12,14,15,16,17].

For space reasons, some proofs and technical details are omitted in this extended abstract, and can be found in the appendix.

2 The Arrow Placement Problem

We first formally define our arrow placement problem and establish its theoretical time complexity. Let $G = (V, E)$ be a digraph and let Γ be a straight-line drawing of G. We assume that in Γ each vertex $v \in V$ is drawn as a circle (possibly a point) C_v. We also assume that, for each edge $e \in E$, the arrow of e is modeled

* Work is partially supported by the MIUR project AMANDA "Algorithmics for MAssive and Networked DAta", prot. 2012C4E3KT_001.

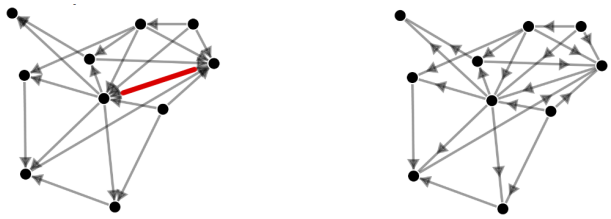

Fig. 1. Layouts of a digraph with 10 vertices and 21 edges. (left) The arrows are placed by a common editor; several arrows overlap and the direction of, e.g., the thick red edge is not clear. (right) The arrows are placed by our exact method.

as a circle C_e of positive radius, centered in a point along the segment that represents e: when Γ is displayed, the arrow of e is drawn as a triangle inscribed in C_e, suitably rotated according to the direction of e. We assume that all circles representing a vertex (arrow) have a common radius r_V (r_E, respectively). We say that two arrows—or an arrow and a vertex—*overlap* if their corresponding circles intersect in two points. An arrow and an edge *overlap* if the segment representing the edge intersects the circle representing the arrow in two points. For the sake of simplicity, we reuse terms of theoretical concepts also for their visual representation: "arrow" and "vertex" also refer to their corresponding circle in Γ; "edge" also refers to its corresponding segment in Γ.

Definition 1. *Let a_e denote the arrow of an edge $e \in E$. A* valid position *for a_e in Γ is such that: (P1) for every vertex $v \in V$, a_e and v do not overlap; (P2) for every edge $g \in E$, $g \neq e$, a_e and g do not overlap. An assignment of a valid position to each arrow is called a* valid placement *of the arrows, denoted by P_Γ.*

Definition 2. *Given a valid placement P_Γ, the* overlap number *of P_Γ is the number of pairs of overlapping arrows, and is denoted as* $\mathrm{ov}(P_\Gamma)$.

Given a straight-line drawing Γ of a digraph $G = (V, E)$, and constants r_V, r_E, we ask for a valid placement P_Γ of the arrows (if one exists) such that $\mathrm{ov}(P_\Gamma)$ is minimum. This optimization problem is NP-hard; we prove this by showing the hardness of the following decision problem Arrow-Placement.

Problem: Arrow-Placement

Instance: $\langle G = (V, E), \Gamma, r_V, r_E \rangle$.

Question: Does there exist a valid placement P_Γ of the arrows with $\mathrm{ov}(P_\Gamma) = 0$?

Theorem 1. *The* Arrow-Placement *problem is NP-hard.*

The proof of Thm. 1 uses a reduction from Planar 3-SAT [11], and is similar to those used in the context of edge labeling [8,12,15,17]. It yields an instance of Arrow-Placement where the search of a valid placement P_Γ with $\mathrm{ov}(P_\Gamma) = 0$ can be restricted to a finite number of valid positions for each arrow. Hence,

Arrow-Placement remains NP-hard even if we fix a finite set of positions for each arrow, and a valid placement with overlap number zero (if any) may only choose from these positions. As this variant of Arrow-Placement, which we call Discrete-Arrow-Placement, clearly belongs to NP, it is NP-complete.

3 Algorithms

We describe algorithms for the optimization version of Discrete-Arrow-Placement. We assume that a set of valid positions for each arrow is given, based on $\{\Gamma, r_V, r_E\}$, and look for a valid placement P_Γ that minimizes $\mathrm{ov}(P_\Gamma)$ over this set of positions. We give both an exact algorithm and two variants of a heuristic, which we experimentally compare in Section 4. Given an edge $e \in E$, let A_e denote the set of valid positions for the arrow of edge e, and let $A := \bigcup_{e \in E} A_e$ be the set of all valid positions. Our algorithms are based on an *arrow conflict graph* C_A, depending on A, Γ, and r_E. The positions A form the node set of C_A. Two positions are *conflicting*, and connected by an (undirected) edge in C_A, if they correspond to positions of different edges and the arrows would overlap when placed on these positions. Finding a valid placement P_Γ with $\mathrm{ov}(P_\Gamma) = 0$ means to select one element from each A_e such that they form an independent set in C_A. More general, finding a valid placement P_Γ with $\mathrm{ov}(P_\Gamma) = k$ ($k \geq 0$) means to select one element from each A_e such that they induce a subgraph with k edges in C_A. Our exact algorithm minimizes k using an ILP formulation, while our heuristic adopts a greedy strategy. Both techniques try to minimize the distance of each arrow from its target vertex as a secondary objective. However, our algorithms can be easily adapted to privilege other positions (e.g., close to the source vertices, in the middle of the edges, etc.), or to consider bidirected edges.

ILP formulation. For each position $p_e \in A_e$ of an edge $e = (v, u)$, we have a binary variable x_{p_e}. We define a distance $d(p_e) \in \{1, \ldots, |A_e|\}$, from p_e to u: $d(p_e) = 1$ ($d(p_e) = |A_e|$) means that p_e is the position closest (farthest, respectively) to u. Let $E_A := E(C_A)$ be the pairs of conflicting positions. For every $(p_e, p_g) \in E_A$, we define a binary variable $y_{p_e p_g}$. The total number of variables is $O(|A|^2)$, and we write:

$$\min \sum_{(p_e, p_g) \in E_A} y_{p_e p_g} + \frac{1}{M} \cdot \sum_{e \in E} \sum_{p_e \in A_e} d(p_e) x_{p_e} \tag{1}$$

$$\sum_{p_e \in A_e} x_{p_e} = 1 \qquad \forall e \in E \tag{2}$$

$$x_{p_e} + x_{p_g} \leq y_{p_e p_g} + 1 \qquad \forall (p_e, p_g) \in E_A \tag{3}$$

The objective function minimizes the overlap number and, secondly, the sum of the distances of the arrows from their target vertices. To do this, the second term is divided for a sufficiently large constant M. For example, one can set $M = |E| \max_{e \in E}\{|A_e|\}$. Equations (2) guarantee that exactly one valid position

per edge is selected. Constraint (3) enforces $y_{p_e p_g} = 1$ if both conflicting positions x_{p_e} and x_{p_g} are chosen. In the following, the exact technique will be referred to as Opt. We remark that optimization problems and ILP formulations similar to above have been given in the context of edge and map labeling [3,4,8,10,14,15].

Heuristics. Our heuristics follow a greedy strategy, again based on C_A. Let $p_e \in A_e \subset V(C_A)$ as above. We initially assigns cost $c(p_e)$ to each position p_e, and then execute $|E|$ iterations. In each iteration, we select a position p_e of minimum cost (over all $e \in E$) and place the arrow of the corresponding edge there; then, we remove all positions A_e from C_A (including p_e), and update the costs of the remaining positions. We define $c(p_e) := \delta(p_e) + \frac{1}{M} d(p_e) + T\sigma_{p_e}$, where: $\delta(p_e)$ is the degree of p_e in C_A (i.e., the number of positions conflicting with p_e); constant M and "distance" $d(p_e)$ are defined as in the ILP; σ_{p_e} is the number of already chosen positions conflicting with p_e (0 in the first iteration); T is equal to the maximum initial cost of a valid position. This cost function guarantees that: (i) positions conflicting with already selected positions are chosen only if necessary; (ii) the algorithm prefers positions with the minimum number of conflicts with the remaining positions and, among them, those closer to the target vertex. Since constructing C_A may be time-consuming in practice (we compare all pairs of valid positions), we also consider using only a subset of the edges of C_A; we may consider only those conflicts arising from positions of adjacent edges in the input graph. In the following, HeurGlobal is the heuristic that considers full C_A, while HeurLocal is the variant based on this simplified version of C_A.

4 Experimental Analysis

Experimental Setting. We use three different sets of graph: Planar are biconnected planar digraphs with edge density 1.5–2.5, randomly generated with the OGDF [2]. Random are digraphs generated with uniform probability distribution with edge density 1.4–1.6. Both sets contain 30 instances each; 6 graphs for each number of vertices $n \in \{100, 200, \ldots, 500\}$. We did not generate denser graphs, as they give rise to cluttered drawings with few valid positions for the arrows—there, the arrow placement problem seems less relevant. Finally, North is a popular set of $1,275$ real-world digraphs with 10–100 vertices and average density 1.4 [13]. We draw each instance of the three sets with straight-line edges using OGDF's FM3 algorithm [5]. The layouts of the Planar may contain edge crossings, as they are generated by a force-directed approach.

Value r_E is chosen as the minimum of (a) 40% of the shortest edge length, (b) 25% of the average edge length, and (c) 10 pixels, but enforced to be at least 3 pixels. We set $r_V := r_E$. For each edge $e = (w, u)$ we compute positions A_e as follows. The i-th position, $i \geq 1$, has its center at distance $r_V + i \cdot r_E$ from target u. We generate positions as long as they have distance at least $r_V + r_E$ from source vertex w. We then remove positions that overlap with edges or vertices in Γ. If no valid positions remain, we choose the one closest to u as e's unique arrow position. Thus, in the final placements there might be some conflicts between an

arrow and a vertex or edge of the drawing. We call such conflicts *crossings* and observe that a single invalid position may result in several crossings.

We apply Opt, HeurGlobal, and HeurLocal to each of the drawings. The algorithms are implemented in C# and run on an Intel Core i7-3630QM notebook with 8 GB RAM under Windows 10. For the ILP we use CPLEX 12.6.1 with default settings. For each computation, we measure total running time, overlap number, and number of crossings (due to invalid positions, see above). From the qualitative point of view, we also compare the algorithmic results with a trivial placement, called Editor, which simply places each arrow close to its target vertex, similarly as most graph editors do. We also measure *placement time*, i.e., the time spent by an algorithm to find a placement *after* C_A has been computed.

Results. For PLANAR, the average numbers of positions in C_A range from 640 to $7,150$. Figs. 2(a) and 2(b) show that for PLANAR all the algorithms are very applicable, although Opt is of course significantly slower. While the pure placement time for HeurGlobal is not much longer than that of HeurLocal, it suffers from the fact that generating the full C_A constitutes roughly 1/3 of its overall runtime, whereas the generation time of the reduced conflict graph is rather neglectable. On the other hand, Fig. 2(c) shows that HeurGlobal practically coincides with the optimum w.r.t. the number of overlaps (its average gap is below 3%; the worst gap is 6.76%). HeurLocal still gives very good solutions, with gaps about half that of Editor. Figure 2(d) shows that our algorithms reduce the number of invalid positions by 33–77% compared to Editor. The number of *crossings* is the same for all our algorithms, as they occur when we cannot find any valid position for arrows during the generation procedure. Figure 2(d) shows that our algorithms cause significantly less crossings than Editor.

For RANDOM, average numbers of positions in C_A range from 640 to $4,377$. The general behavior for RANDOM is similar to that of PLANAR but the difference between the running time of Opt and the heuristics is slightly more pronounced (Fig. 3(a)). Again, constructing C_A constitutes roughly 1/3 of HeurGlobal's running time. Still, the quality of HeurGlobal's solutions again essentially coincide with Opt; the other heuristics are now closer than before, see Fig. 3(b).

For NORTH, the average $|V(C_A)|$ range from 62 to 311. We observe the same patterns, see Figs. 4: HeurLocal requires nearly no time, while HeurGlobal is very competitive at just above 20ms for the large graphs (a third of which is the construction of full C_A). Again, Opt always finds a solution very quickly, in fact within roughly 80ms. HeurGlobal again gives essentially optimal solutions, while HeurLocal exhibits 5–10% gaps. Editor requires 30–50% more overlaps than Opt.

5 Conclusions and Future Work

We discussed optimizing arrow head placement in directed graph drawings, to improve readability. As mentioned, this is very related to studies in map and graph labeling, but its specifics seem to make a more focused study worthwhile.

Our techniques are of practical use, and could be sped-up by constructing C_A using a sweepline or the labeling techniques in [16]. It would be interesting

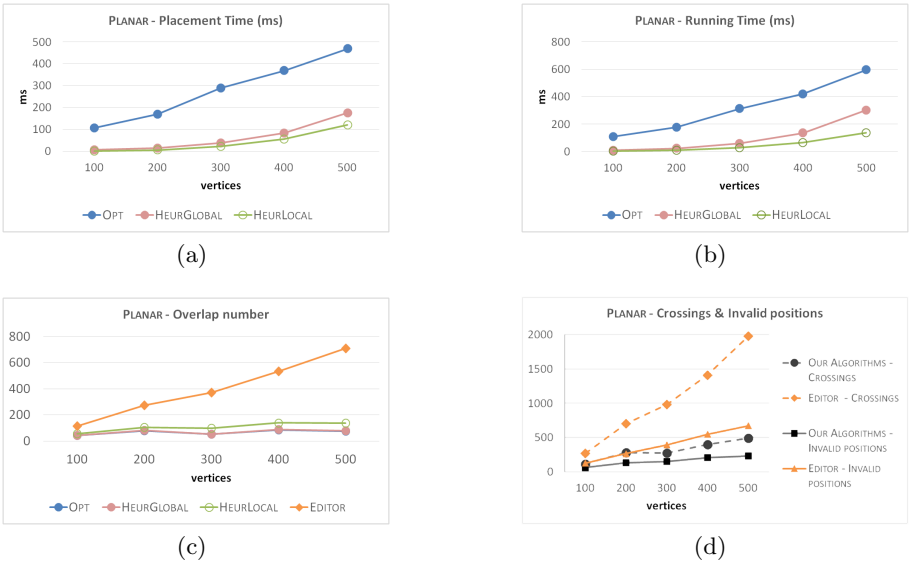

Fig. 2. PLANAR: (a) placement time; (b) total running time; (c) number of overlaps; (d) number of crossings (edge/vertex with arrow) and of invalid positions.

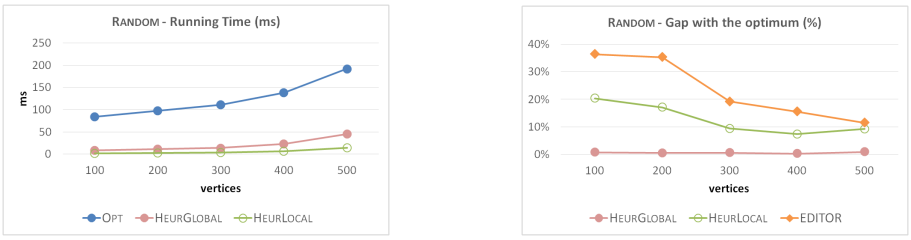

Fig. 3. RANDOM: (a) total running time; (b) number of overlaps, relative to Opt.

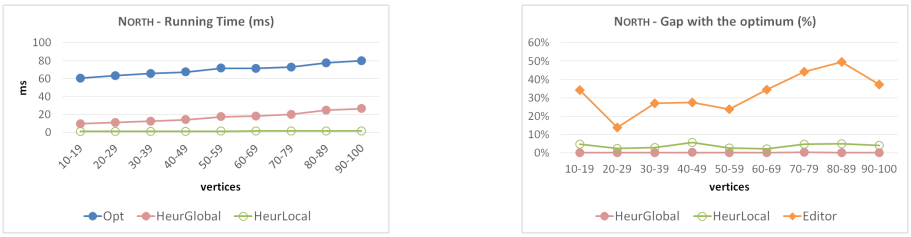

Fig. 4. NORTH: (a) total running time; (b) number of overlaps, relative to to Opt.

to validate the effectiveness of our approach through a user study (e.g. for tasks that involve path recognition). Moreover, one may consider both placing labels and arrow heads. Finally, the non-discretized problem variant, as well as the variants' respective (practical) benefits, should be investigated in more depth.

Acknowledgments. Research on this problem started at the Dagstuhl seminar 15052 [1]. We thank Michael Kaufmann and Dorothea Wagner for valuable discussions, and the anonymous referees for their comments and suggestions.

References

1. Brandes, U., Finocchi, I., Nöllenburg, M., Quigley, A.: Empirical Evaluation for Graph Drawing (Dagstuhl Seminar 15052). Dagstuhl Reports 5(1), 243–258 (2015)
2. Chimani, M., Gutwenger, C., Jünger, M., Klau, G.W., Klein, K., Mutzel, P.: The open graph drawing framework (OGDF). In: Tamassia, R. (ed.) Handbook of Graph Drawing and Visualization, chap. 17. CRC Press (2014), `www.ogdf.net`
3. Gemsa, A., Niedermann, B., Nöllenburg, M.: Trajectory-based dynamic map labeling. In: Cai, L., Cheng, S., Lam, T.W. (eds.) ISAAC 2013. LNCS, vol. 8283, pp. 413–423. Springer (2013)
4. Gemsa, A., Nöllenburg, M., Rutter, I.: Evaluation of labeling strategies for rotating maps. In: Gudmundsson, J., Katajainen, J. (eds.) SEA 2014. LNCS, vol. 8504, pp. 235–246. Springer (2014)
5. Hachul, S., Jünger, M.: Drawing large graphs with a potential-field-based multilevel algorithm. In: Pach, J. (ed.) GD 2004. LNCS, vol. 3383, pp. 285–295. Springer (2004)
6. Holten, D., Isenberg, P., van Wijk, J.J., Fekete, J.: An extended evaluation of the readability of tapered, animated, and textured directed-edge representations in node-link graphs. In: IEEE PacificVis 2011. pp. 195–202. IEEE (2011)
7. Holten, D., van Wijk, J.J.: A user study on visualizing directed edges in graphs. In: CHI 2009. pp. 2299–2308. ACM (2009)
8. Kakoulis, K.G., Tollis, I.G.: On the complexity of the edge label placement problem. Comput. Geom. 18(1), 1–17 (2001)
9. Kakoulis, K.G., Tollis, I.G.: Labeling algorithms. In: Tamassia, R. (ed.) Handbook on Graph Drawing and Visualization., pp. 489–515. Chapman and Hall/CRC (2013)
10. van Kreveld, M.J., Strijk, T., Wolff, A.: Point labeling with sliding labels. Comput. Geom. 13(1), 21–47 (1999)
11. Lichtenstein, D.: Planar formulae and their uses. SIAM J. Comput. 11(2), 329–343 (1982)
12. Marks, J., Shieber, S.: The computational complexity of cartographic label placement. Tech. rep., Technical Report 05-91, Harvard University (1991)
13. North graphs. `http://www.graphdrawing.org/data.html`
14. Strijk, T., van Kreveld, M.J.: Practical extensions of point labeling in the slider model. GeoInformatica 6(2), 181–197 (2002)
15. Strijk, T., Wolff, A.: Labeling points with circles. Int. J. Comput. Geometry Appl. 11(2), 181–195 (2001)
16. Wagner, F., Wolff, A., Kapoor, V., Strijk, T.: Three rules suffice for good label placement. Algorithmica 30(2), 334–349 (2001)

17. Wolff, A.: A simple proof for the np-hardness of edge labeling. Tech. rep., Technical Report 11/2000, Institute of Mathematics and Computer Science, Ernst Moritz Arndt University Greifswald (2000)

Appendix A – Additional examples

Figure 5 depicts some larger real-world layouts from our graph benchmark. Examples of layouts for larger artificial instances are shown in Fig. 6. When the visual complexity of the graph layout increases, the arrow placement problem seems to become less relevant: indeed, cluttered drawings have few valid positions for the arrows and also the solutions of our algorithms may contain several crossings between arrows and other objects; see, e.g., Figs. 6(c) and 6(d).

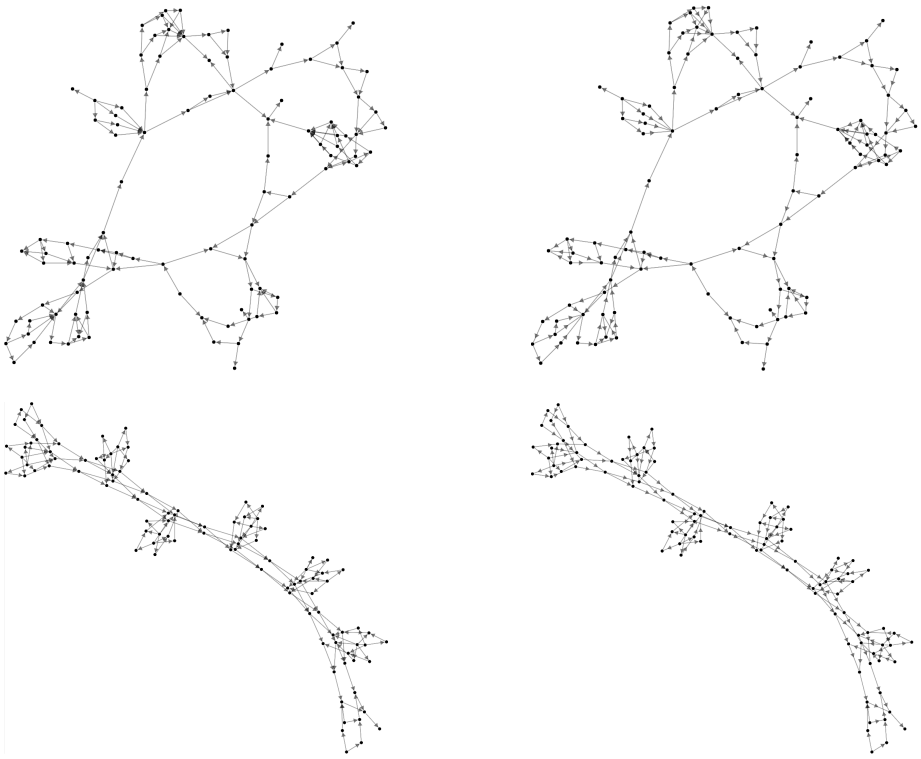

Fig. 5. Drawings of two North digraphs with 99 vertices and 148 edges. On the left, arrows are placed by a common editor; several arrows overlap. On the right, arrows are placed by our exact technique.

Appendix B – Additional Charts from the Experiments

We report some additional charts from the experiments of Section 4; see Figs. 7 and 8. Moreover, in order to evaluate the scalability of our approach, we extended both Planar and Random sets with 30 larger instances each (6 graphs

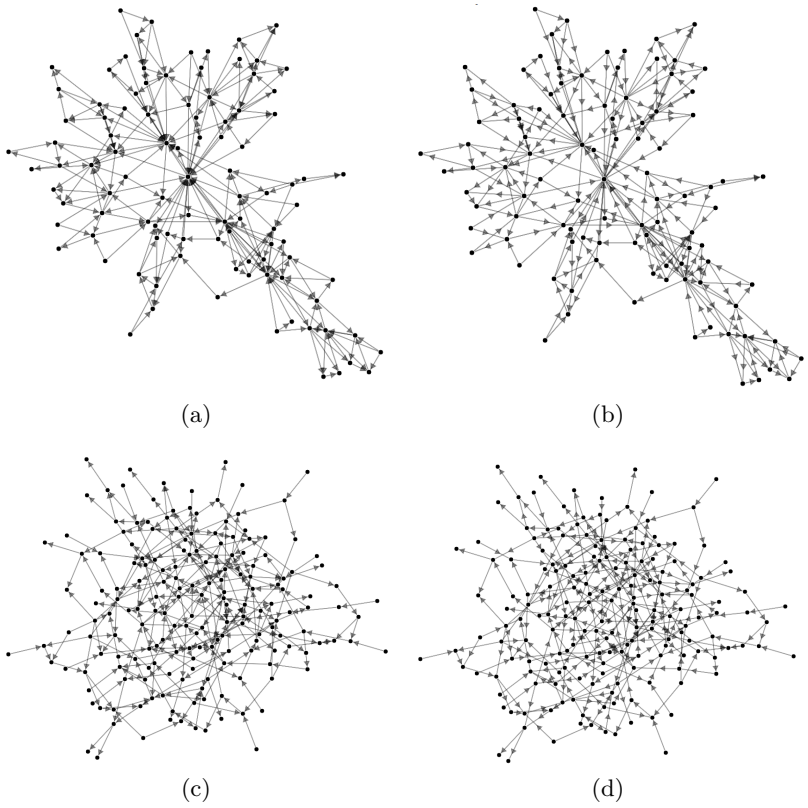

Fig. 6. (a)&(b) Drawings of a PLANAR digraph with 100 vertices and 250 edges; (c)&(d) Drawings of a RANDOM digraph with 200 vertices and 280 edges. In the drawings on the left side the arrows are placed by a common editor; in the drawings on the right side the arrows are placed by our exact technique.

for each number of vertices $n \in \{600, 700, \ldots, 1000\}$) and ran our algorithms on them. The behavior of the algorithms is similar to that reported for smaller instances: the quality of HeurGlobal's solutions almost coincides with the quality of Opt's solutions; see Figs. 9(c), 10(c). Constructing C_A remains the most expensive procedure, especially for the larger instances in the PLANAR set; see Figs. 9(a), 9(b). Finally, our algorithms still generate significatively less crossings than Editor; see Figs. 9(e), 10(e). Since Editor can use positions that are not valid for our algorithms, its performance in terms of number of overlaps tends to compare with the performances of our heuristics on some of the most cluttered layouts (see, e.g., the largest instances of RANDOM).

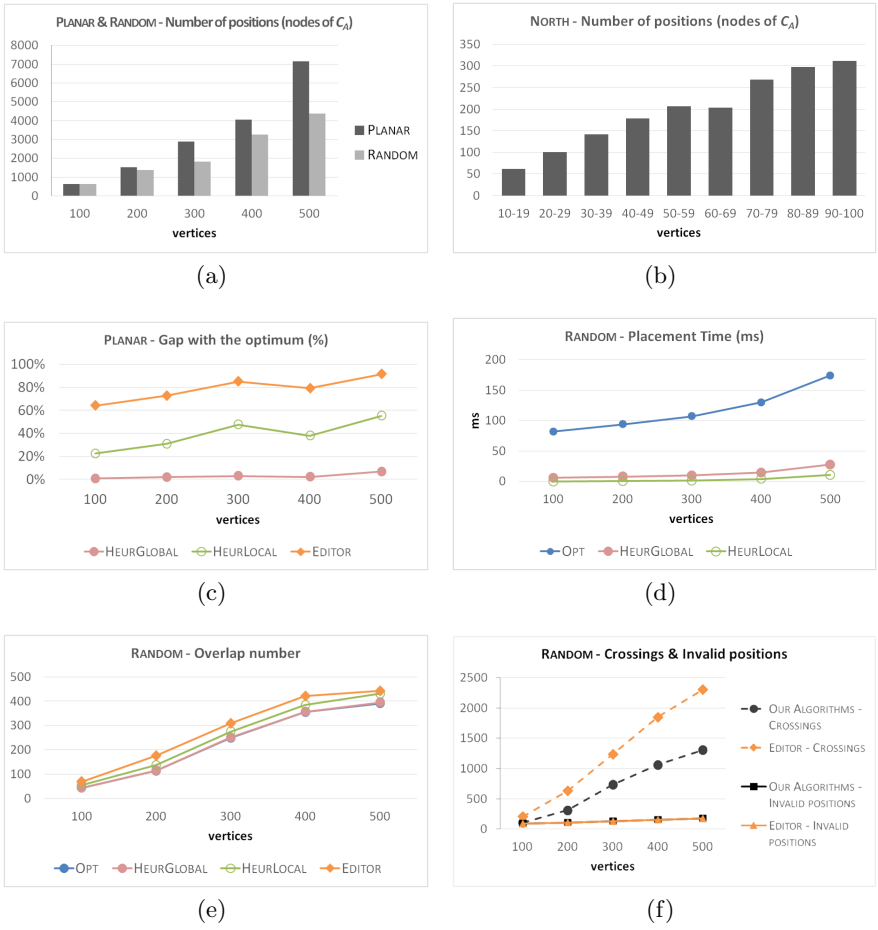

Fig. 7. Number of positions (nodes in C_A) for: PLANAR and RANDOM (a), and NORTH (b). — PLANAR: (c) number of overlaps, relative to Opt. — RANDOM: (d) placement time; (e) number of overlaps; (f) number of crossings (between an edge/vertex and an arrow) and of invalid positions.

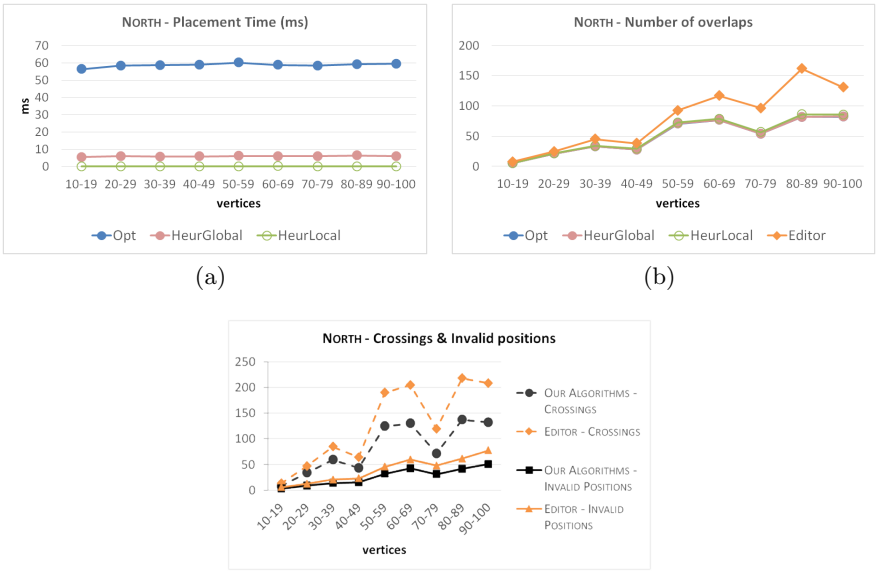

(a) (b)

(c)

Fig. 8. North: (a) placement time; (b) number of overlaps; (c) number of crossings (between an edge/vertex and an arrow) and of invalid positions.

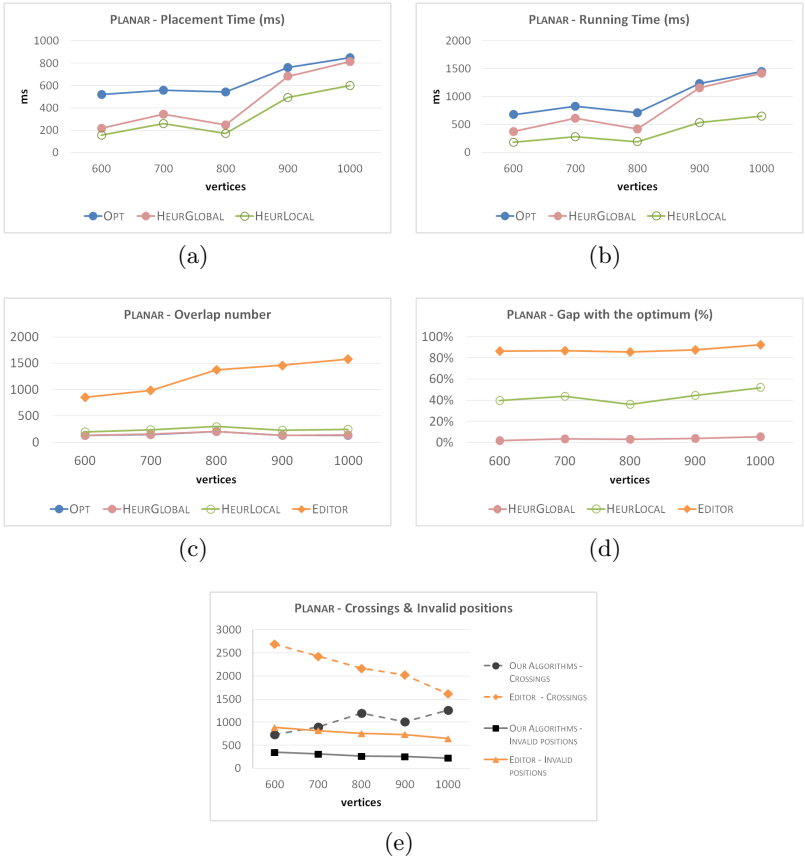

Fig. 9. Larger graphs of PLANAR: (a) placement time; (b) total running time; (c) number of overlaps; (d) number of overlaps, relative to Opt; (e) number of crossings (between an edge/vertex and an arrow) and of invalid positions.

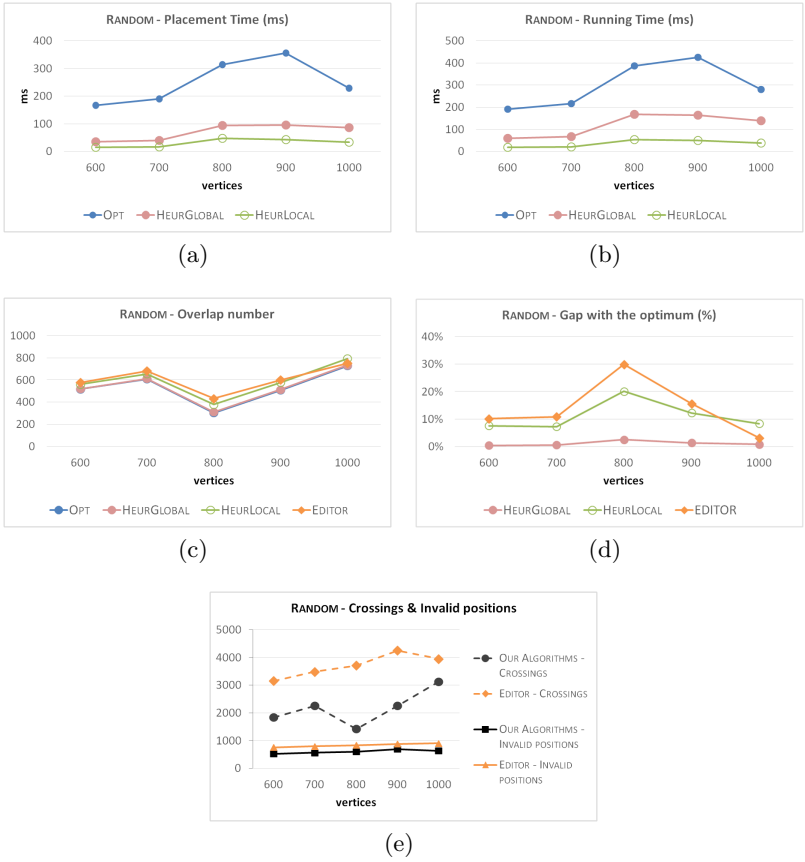

Fig. 10. Larger graphs of RANDOM: (a) placement time; (b) total running time; (c) number of overlaps; (d) number of overlaps, relative to Opt; (e) number of crossings (between an edge/vertex and an arrow) and of invalid positions.

Appendix C – NP-hardness Proof

Theorem 1. *The* Arrow-Placement *problem is NP-hard.*

Proof. We show that Arrow-Placement is NP-hard by a reduction from Planar 3-SAT (P3SAT) [11], the special case of 3-SAT in which the bipartite graph of variables and clauses is planar. Our proof follows the proof of Knuth and Raghunathan [3] for the Metafont labeling problem in which the variables are placed onto a horizontal line, while the three-legged clauses are placed above or below them. The clauses are properly nested so that none of the legs between clauses and variables cross each other. The connections can either be positive or negative. We remark that our reduction technique is similar to those used in the context of edge and point labeling: labeling points with circular labels [15] and with axis-parallel rectangular labels in different models [12], [10]; labeling edges with axis-parallel rectangular labels [17]. To reduce P3SAT to Arrow-Placement, we describe how to construct from any planar 3-CNF formula a (possibly) connected drawing Γ such that there exists a valid placement P_Γ of the arrows with $\mathrm{ov}(P_\Gamma) = 0$ if and only if the formula is satisfiable. Let ϕ be an instance of the P3SAT problem consisting of the set of clauses $c_1, \ldots, c_m$, each having three literals from the Boolean variables $x_1, \ldots, x_n$. We now explain how to construct a variable gadget for each variable x_i and a clause gadget for each clause c_i of ϕ. We use two building blocks.

The first building block, which we call *triangle-block*, is shown in Fig. 11(a): it is composed of an equilateral triangle T having side length equal to $2r_E\sqrt{3}$. By basic geometry, it follows that along each edge e of T a circle C_e (modeling arrow a_e) with its center on e and radius r_E admits only two alternative valid positions represented in the figure by circles with solid and dashed boundaries, respectively. In the following, we simply call a valid position represented by a circle with solid (dashed) boundary a *solid (dashed) position*. It can be observed that there exist only two valid placements for T, one called *solid placement* that uses only solid positions and one called *dashed placement* that uses only dashed positions. We call *base* of T, the horizontal edge of T, while we call the other two edges of T the *left* and the *right* edges of T, respectively.

The second building block which we call *trapezoid-block* is a right trapezoid Q in which the vertical side is missing. We call the two horizontal edges of Q the *major base* and the *minor base* of Q and the third edge of Q the *diagonal side* of Q. As shown in Fig. 11(b), we have two variants of Q (called Q-*left* and Q-*right*) that are symmetric. The lengths of the three edges of Q are chosen as follows. Denote by q_1 and q_2 (q_3 and q_5) the two end-points of the minor (major) base of Q and denote by q_4 the projection of q_1 along the major base. Note that $\overline{q_1q_4}$ is the *height* of Q. We have the following relations: $\overline{q_3q_4} = \overline{q_1q_4}$, $\overline{q_1q_2} = \overline{q_4q_5}$. The minor base of Q has length $2r_E\sqrt{3}$, the height of Q has length $r_E(1 + \frac{\sqrt{2}}{2})$, then the major base of Q has length $r_E(1 + \frac{\sqrt{2}}{2}) + 2r_E\sqrt{3}$. Observe that the

[3] Knuth, D.E., Raghunathan, A.: The problem of compatible representatives. SIAM J. Discrete Math. 5(3), 422-427 (1992)

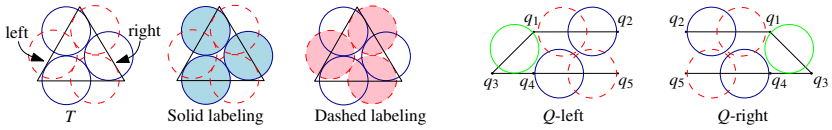

Fig. 11. (a) A triangle-block T and the two possible valid placements of T: the solid placement and the dashed placement. (b) The two variants Q-left and Q-right of a trapezoid-block Q. The diagonal side admits a unique valid position, that is represented by a green circle.

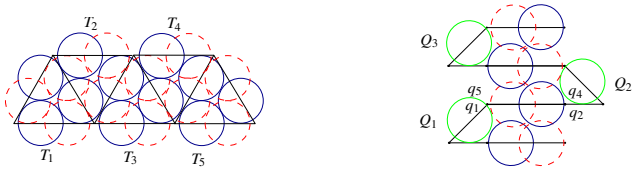

Fig. 12. (a) A variable gadget $\mathcal{T} = \{T_1, \ldots, T_5\}$. (b) A chain $\mathcal{Q} = \{Q_1, Q_2, Q_3\}$ of trapezoid-blocks, which represents the leg of a clause gadget. As example, the two end-points q_1 and q_2 of the minor base of Q_1 coincide with the points q_5 and q_4 of the major base of Q_2.

height of Q is less than $2r_E$ and that the diagonal side of Q is the diagonal of a square with side $\overline{q_3q_4}$. By basic geometry it follows that: (i) the diagonal side of Q admits a unique valid position that is tangent to segment $\overline{q_3q_4}$; (ii) the minor and major bases of Q admit only two alternative valid positions. Also in this case it can be observed that beside the unique valid position of the diagonal side, the selected valid positions of the minor and major bases of Q must be either both solid or both dashed. A valid placement of Q using the unique valid position of the diagonal side and the two solid (dashed) positions of the minor and major bases is called a *solid (dashed) placement*.

Variable gadget. The variable gadget is composed by a horizontal chain $\mathcal{T} = \{T_1, \ldots, T_k\}$ of triangle-blocks, for a given odd $k \geq 5$. Each triangle block T_i with i even is an upside-down triangle that shares its left edge with the right edge of the triangle T_{i-1} and its right edge with the left edge of the triangle T_{i+1} as shown in Fig. 12(a). This implies that the choice of one between the solid placement or the dashed placement for the leftmost triangle T_1 enforces the same kind of placement for all other triangle-blocks of the chain. Hence, we call a *solid (dashed) placement* of $\mathcal{T}$, a placement of $\mathcal{T}$ composed of only solid (dashed) positions. In our construction, the truth (false) value of a variable x_i is encoded by a solid (dashed) placement of the corresponding variable gadget $\mathcal{T}^i$.

Clause gadget. The clause gadget is composed of three *vertical legs* and one *horizontal part* (refer to Fig. 13). The horizontal part is composed of two hor-

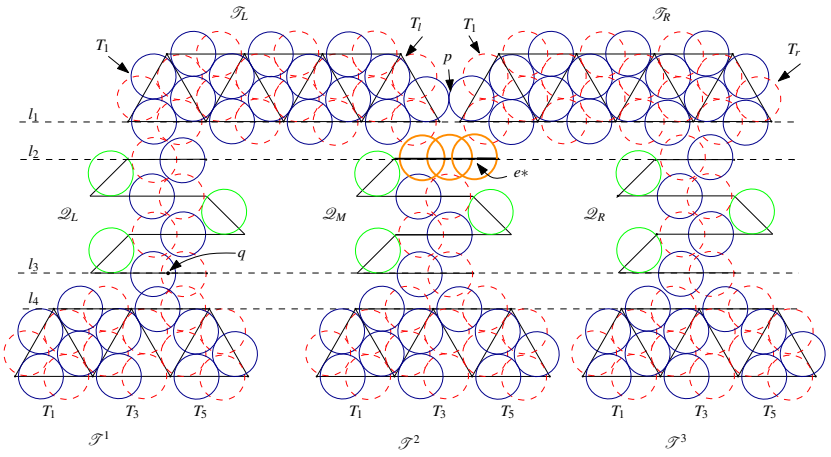

Fig. 13. Illustration of a clause gadget. The clause of a boolean formula ϕ modeled by this example is $c_1 = (\overline{x_1} \vee x_2 \vee x_3)$. Note that if x_1 is true and x_2 and x_3 are false, then edge e^* does not admit a valid position.

izontal chains $\mathcal{T}_L = \{T_1, \ldots, T_l\}$ and $\mathcal{T}_R = \{T_1, \ldots, T_r\}$ of triangle-blocks (see Fig. 13). We have the following properties: (i) The solid and dashed positions along the edges of the triangle-blocks of $\mathcal{T}_R$ are inverted with respect to those of the triangle-blocks of $\mathcal{T}_L$; (ii) For every odd i, the base of $T_i \in \mathcal{T}_L$ and of $T_i \in \mathcal{T}_R$ lie on the same horizontal line l_1; (iii) The solid position of the right edge of $T_l \in \mathcal{T}_L$ is tangent to the solid position of the left edge of $T_r \in \mathcal{T}_R$. Denote by p this tangent point.

We now describe the three legs, called *left leg*, *middle leg*, and *right leg*. Each leg is a vertical chain $\mathcal{Q} = \{Q_1, \ldots, Q_k\}$ (with k odd) of trapezoid-blocks composed in the following way (see Fig. 12(b)): each Q_i with i even is a Q-right trapezoid-block that shares its major base with the minor base of Q_{i-1} (that is a Q-left trapezoid). We denote by $\mathcal{Q}_L$, $\mathcal{Q}_M$ and $\mathcal{Q}_R$ the vertical chains of the left, the middle and the right legs, respectively. We denote by the *top base (bottom base)* of each leg the minor base (major base) of the topmost (bottom most) trapezoid-block Q_k (Q_1) of each leg. We have the following properties: (i) $\mathcal{Q}_L$, $\mathcal{Q}_M$, $\mathcal{Q}_R$ have the same number k of trapezoid-blocks, and therefore, their top bases (bottom bases) lie on the same horizontal line l_2 (l_3); (ii) The top base of $\mathcal{Q}_M$ is longer than those of $\mathcal{Q}_L$ and $\mathcal{Q}_R$ so that this edge (denoted by $e*$) is the only one that admits more than two valid positions. We set the length of $e*$ equal to $r_E(4\sqrt{3} - 2)$.

We place the legs below the horizontal part in the following way: we set the Euclidean distance d between l_1 and l_2 equal to $r_E\left(1 + \frac{\sqrt{2}}{2}\right)$, that is $d < 2r_E$ and for the left leg, we vertically align the middle point of the top base of $\mathcal{Q}_L$

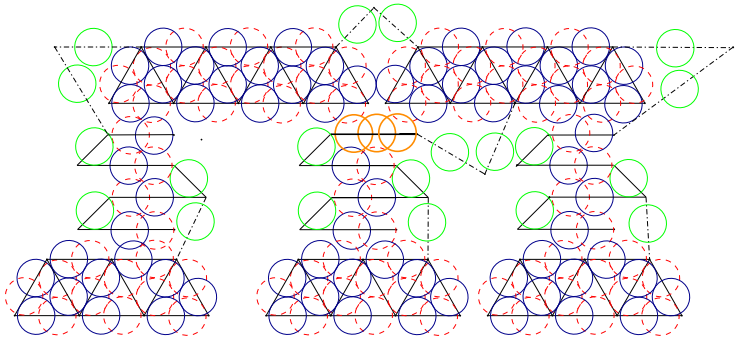

Fig. 14. The clause of Fig. 13 where some other edges have been added to make Γ connected. The added edges are represented by dashed dotted edges and admit always a valid position.

with the middle point of the base of $T_1 \in \mathcal{T}_L$; for the middle leg we vertically align the middle point of the top base of $\mathcal{Q}_M$ with point p; for the right leg we vertically align the middle point of the top base of $\mathcal{Q}_R$ with the left end-point of the base of $T_r \in \mathcal{T}_R$.

The above geometric relations guarantee the following properties: **(a)** The choice of one between the solid position and the dashed position for the bottom-base of $\mathcal{Q}_L$ ($\mathcal{Q}_R$) enforces the same kind of positions for all the other trapezoid-blocks of the leg and for all the triangle-blocks of $\mathcal{T}_L$ ($\mathcal{T}_R$). The same holds for the trapezoid-blocks of $\mathcal{Q}_M$. That is, the truth or the false value is transmitted along each leg and along the horizontal part of the clause. **(b)** If each leg transmits the false value (i.e. the enforced positions are dashed), then the length of edge $e*$ guarantees that there does not exist a valid position for $e*$. Otherwise, if at least one of the three legs transmits the truth value (i.e. at least one leg uses solid positions), then there exists a valid position for edge $e*$.

In Fig. 13 we also show how each clause leg is connected to a variable gadget. We align the bases of the upside-down triangles of each variable gadget $\mathcal{T}^i$ (corresponding to a variable x_i) along a horizontal line l_4 and, as in the case of l_1 and l_2, we set the Euclidean distance between l_3 and l_4 equal to $d = r_E\left(1 + \frac{\sqrt{2}}{2}\right)$. Depending on the literal we attach the clause leg to the corresponding variable gadget in one of two possible places. Denote by q the middle point between the centers of the solid and dashed positions of the bottom base of a leg. If the literal is negated we vertically align q with the middle point of the base of a triangle-block T_i with i even (i.e. an upside-down triangle) of the variable gadget $\mathcal{T}^i$. If the literal is not negated, we vertically align q with the end-point shared by the bases of the two triangle blocks T_i and T_{i+2} with i even.

If all literals of a clause are false, then the edge $e*$ does not admit a valid position. Indeed all the three legs of the clause transmits the false value (i.e. the selected positions of all the three legs are dashed). Otherwise if at least one literal is true, then $e*$ and all the other edges admit a valid position. Indeed in this case, the clause leg attached to the true literal transmits the truth value (i.e. the selected positions for it are solid). Therefore, given any planar 3-CNF formula ϕ there exists a valid placement P_Γ of the arrows such that $\mathrm{ov}(P_\Gamma) = 0$ if and only if ϕ is satisfiable.

Moreover the construction guarantees that the variables are placed onto a horizontal line and they can be extended to reach all the clauses. The same holds for the three legs and the two horizontal parts of the clauses.

Observe that the drawing Γ described so far is not connected. In order to make Γ connected, we can add a suitable number of extra edges, represented in Fig. 14 as dashed-dotted segments, that always admit a valid position. □